



University of Auckland

Department of Computer Science

Harmony on an
Expanding Net

Barry Fenn and Hermann Maurer

Auckland
New Zealand

Submitted, by request, for inclusion
in the October issue of interactions -
the new ACM journal of CHI activities.

Harmony on an Expanding Net

Barry Fenn and Hermann Maurer

Harmony on an Expanding Net

Barry Fenn and Hermann Maurer

The Internet is experiencing an exceptional rate of growth, based upon the systems for its use having reached a critical level of usability. The development of 'second generation' network systems, such as Hyper-G and its Harmony viewer, can now boost this activity past another threshold of capability.

Overview

Recent public-relations hype from the telecommunications sector and other highly-centralised media enterprises paints a future of video-on-demand, 500-channel television, and vastly extended opportunities for interactive home shopping. This is often described as a forthcoming 'Information Superhighway'... but the more appropriate real-world model is a 'Railroad': This relates to monopoly ownership of the separate transport systems, local hardware standards, and quorum-determined scheduling for efficient operation.

There's a growing consensus, however, among those who are looking most thoughtfully at what's happening, that such visions of forthcoming centrally-oriented multimedia-delivery empires are going to be left in the digital dust of the much more feral developments already taking place around the **Internet** and associated subscription networks. A very diverse 'Virtual Community' is beginning to form which bears much more similarity to that road transport metaphor; right down to electronic country lanes leading to rural telecottages.

Increasing numbers of individuals, educational bodies, and businesses of all sizes are being attracted to this existing 'network of networks' as a means of relating directly with each other, rather than having to depend on the services of a few principal mass-media purveyors. Not that centralised services and interactive resource 'databases' won't be readily available through whatever the Internet evolves into during the next decade. Just that the real driving force of the whole thing is clearly going to be a range of essentially **new modes of communication between people**.


Barry Fenn (barry@cs.auckland.ac.nz) is the External Coordinator for the University of Auckland's HyperMedia Unit and a Computer/Human Interaction specialist.

Professor Hermann Maurer (hmaurer@iicm.tu-graz.ac.at) heads the Institute at Graz University of Technology which initiated development of Hyper-G and Harmony and is also the founder of the University of Auckland's HyperMedia Unit.

What we now know as the Internet has evolved from the US Defence Department's ARPANet, a computer communications network that was specifically designed to survive a nuclear war - through having a working structure that is inherently decentralised. In short, it was made to be as autonomous and 'hard to kill' as possible. This military network was first extended to the American university research community, with increasing levels of support from the NSF, then naturally spread to other countries and is now rapidly extending out into the commercial sector. Due to its dispersed nature, it is difficult to even keep track of exactly what's happening on 'the Net'.

The nonprofit **Internet Society** has now assumed responsibility for setting the 'rules of the road' for Internet traffic. They provide a much needed forum for maintaining and evolving basic standards, but take no direct part in defining what sorts of things are actually happening on the Internet, nor the kinds of interactive tools available to its users. The Net has taken on a life of its own and the overall pattern of its future development is *effectively out of the control of any individual government or organisation*.

At the moment there are perhaps 50 million people using the Internet on a regular basis, but the current growth rate is about 15% per month (!) and this could well continue until almost all of those in the 'developed world' are connected (*Figure 1*). A HUGE social and economic momentum exists to drive such an expansion. The telling point is that if this same rate of increase were able to be sustained world wide, everyone within our expanding human population would be connected just after the turn of the millennium, say by 2002. Is anything like this even remotely conceivable?



Putting aside the technical and economic issues involved, this implies a monumental amount of new interactive design work; bringing twenty-first century connectivity to all those whose prior experience with computers has left them justifiably wary - and also to a wide range of people who have only recently emerged from local isolation into third-world poverty. Just imagine, though, what the world would be like if we can somehow manage to get our act *that* together. Collaboration through the Internet itself could provide the means to engineer such a transformation.²

The exponential nature of the Internet's growth bears a direct relationship to some of the biggest issues confronting us here on the planet (*Figure 2*). Will electronic networking simply increase the gulf between the 'have' and the 'have not' sectors within human society, or instead serve as the critical means by which we begin to effectively bridge the existing economic disparities and cultural tensions? Will we find ourselves within a potentially *fatal polarisation* or a truly '*Global Village*'?

The many intertwined patterns of exponential growth going on within our closed system suggest no stable middle ground between the two extremes.




Figure 2

There are no precedents within past human history by which to predict the overall results from changes of this magnitude.³

E-mail and direct file transfers are still the most common form of Internet traffic, but more interactive distributed applications are the real growth sectors. In the absence of past accountability, some universities have recently discovered that up to half of their Internet bill has been going toward *MUD*; short for Multi-User Dungeon role-playing games, in which a number of participants from around the world join each other online for lengthy sessions of interactive competition. Keeping in mind the fact that this particular activity is mostly on the part of a technically-biased selection of mainly male players, we believe that this still says something important about the likely future of the Internet. As Brenda Laurel expressed it; what's happening is the opportunity, "...to be in the presence of a whole bunch of other human beings, any time day or night, any place in the world. That is the fundamentally compelling thing about network communications."⁴

What happens, though, once there are many thousands of times more 'venues' in which to meet and - in particular - how will we manage to find any desired resource once we have most of recorded human culture and knowledge on hand to share with each other?

Being able to present high-quality media to each other and use virtual simulations as a natural extension of everyday communication has a tremendous appeal. Of course we'll need to keep track of all those we relate with, to know what version of any material or application we're dealing with, and to easily find a way to any relevant new resources. Yet we still want the experience to be much like that of being in the middle of a relaxed conversation, turning around to select a book off one's own shelf, and then opening it to a marked page from which to present something - rather than like spending time alone in a library searching for material that may not even exist.

The underlying structure of a distributed network environment needs to be designed from the start with this sort of media-rich communication in mind, in order to elegantly meet such requirements.

The Changing State-of-the-Art on the Net


There have been a series of increasingly capable 'global' systems developed for organising access to resources on large distributed databases. Four in particular are or are in the process of being implemented on a significant scale within the Internet:

Gopher, WAIS, World Wide Web, and Hyper-G.

All four are based on a client/server architecture, whereby the end-user employs client viewer software designed for their own computer platform (and sometimes even adjusted to their own individual preferences) as a means of relating to servers which all use their network environment's standard protocol and information structure. In general the different systems provide open access to the other's clients, so that each viewer can interface with the other's server sites - but only up to the limits of its own native capability or the level of hypermedia support available from that 'alien' server.


Gopher is the oldest, most widespread, and least interactively friendly of these network environments. It is based on a menu-like interface metaphor which allows for keyword searching, but no hypertext linkage (*Figure 3*). You have to know the general sort of thing you're looking for and the location of the particular place where you might find it. Extensions of the basic Gopher protocol have proved quite useful within limited contexts like Campus-Wide Information Systems, but they reach fundamental limits when applied to a much larger and more diverse database like the Internet. You don't become lost or disoriented only because you can only go off to one known place at a time - and then just to retrieve some content.

Figure 3




WAIS (from Wide Area Information System) is best known for its powerful full-text search facility, which allows complex search routines based on natural language requests (*Figure 4*). This has found particular favour with the US government and various publishing companies who want to provide support for those seeking specific content from within a single large database. No hypertext capability or communication support is available, nor does WAIS provide any means to visualise the structure of its information content. Its most important feature is probably the capability for *interoperation*: The other open network environments are thus able to support the WAIS search engine for use with their own contents.

Figure 4



World Wide Web (also referred to as WWW, W³, or ‘the Web’) is the fastest growing domain of the Internet. This success is generally acknowledged as being the result of its separately-developed *Mosaic* viewer (*Figure 5*). The Mosaic interface allows point-and-click access to hypertext links, keyword searching, and rudimentary modes of orientation within the ‘hyperspace’ of links established between individual documents (which can be either text and pictures, a sound recording, or a video sequence).

Figure 5



The design of the World Wide Web limits interaction to a fairly linear pattern of usage; hyperlinks held internal to each document can be followed within each server and off to other sites - or are later used from system memory to backtrack along the course previously taken. The hypertext model on which the Web was originally based presumed that a single user would just want to repeatedly go away from an existing core document, in order to explore a few related resources, and then return to carry on from that prior departure point. Wider-ranging forays, where one ‘surfs’ across content dispersed throughout the Internet, can quickly lead to the light-headed disorientation of being “lost in hyperspace”. No native support is available for annotation of documents or communication between users.

The Mosaic interface, though, has clearly demonstrated a critical threshold of general usability and offers a tantalising taste of a fully networked world’s flavour. It is friendly enough, and permits a useful enough range of interactivity, to have drawn a flood of new users to the Internet. Commercial interest in the development of Mosaic-based services is running high and a subculture of dedicated ‘mouse potatoes’ has even begun to emerge (the name being a take-off on television-centered ‘couch potatoes’).

The constraints imposed by the underlying World Wide Web network environment are also becoming apparent, however. It is common practice for organisations setting up a Web server site to install essentially duplicate systems for either internal or public access, with the separation between them serving as a 'firewall' for security purposes. This is necessary to counteract the anonymous nature of the environment. Work is underway to retrofit a range of extra applications on top of the existing Web architecture - allowing for full-text searching, authentication, user identification, transactional charging, and so on - but this additional functionality comes at the price of lost interface uniformity; different Web servers start to behave differently. In any case, each Web site will always remain a fundamentally separate interactive context, only potentially connected to other specific locations on the Internet through its own cloistered hyperlinks.

Hyper-G is the most recent large-scale networking project and was designed from the start to transcend the limitations of the other existing environments. In order to facilitate this, a closer relationship is maintained between a client and its designated 'home' server (which may both reside within the same workstation, but can also effectively be located miles apart on a Wide Area Network). The most fundamental difference in network architecture is that all of the hypermedia connections are stored and maintained separately from the content, through a parallel *link server* system. This allows material within an arbitrary number of distributed content servers, or even the entire global Hyper-G database, to be experienced as a seamless whole.

Rather than having to follow links from server to server, seeking one site that at least contains a direct reference to some desired resource, the scope of a single Hyper-G search can be adjusted so that it is conducted across any chosen section of the entire network. A complex searching strategy can be initiated which progresses, with increasing refinement, within the regions highlighted by each previous component of the routine.⁵ For instance, one could ask for *all* resources on *Whales* which contain keyworded material about *Orca* as the first part of a fuzzy full-text search for information on their *communication patterns*, specifically as they relate to *kinship bonds* between *female siblings* - thus accomplishing a series of prefiltered searches that would efficiently winnow through a total world-wide database.

The separation between content and links within the Hyper-G environment also facilitates the sort of personal annotation or 'harmless' modification of documents that Ted Nelson envisioned for his full Xanadu system.⁶ This is supported by a graduated series of ways in which users can be recognised by the system; from *Anonymous*, to *Anonymously Identified* (known to the system and others only by a passworded pseudonym), to *Semi-identified* (fully known to the system but appearing to others under a pseudonym), through to fully *Identified*. The author of any document (or 'personalised' version thereof) has the opportunity to specify the details of who may access, publicly annotate, or even directly change their arrangement of material.

The capacity within Hyper-G to positively identify users and limit access privileges means that less elaborate firewalls can provide high levels of security. It also allows the same network environment to be used as a direct means of communication between individual users. The development of software support for this is planned, with an integrated Internet

e-mail protocol being the immediate goal. Direct online interaction, allowing multiple users to share the same hypermedia environment, is included within the overall specification and will be implemented within a distance-learning research project that is just getting underway.

The 'documents' themselves can be any form of media, including 3-D visualisations, and are referenced by Hyper-G in three basic ways; *Collections*, *Clusters* and *Tours*. *Collections* are the sort of hierarchical menu structures employed within the other network environments. *Clusters* of documents can also be arranged, so that they all open and function together in a synchronised way. This allows sound to accompany video, or other truly multimedia combinations. One particular application of clustering is the creation of 'annotated movies', wherein text or other relevant material is linked to a core video presentation and made interactively accessible.

A special feature of Hyper-G is the capability of supporting *Tours*. These are a scripted sequence of links, activated between a series of documents or clusters, which the user may just sit back and watch like a television program. The content is composed of actual network resources, however, so that one can stop at any point to explore some material in more depth or follow any other pattern of hypermedia linkage. Thus an 'article' like this could be made available in the form of synchronised spoken material and text, with the 'illustrations' coming from real online resources scattered about the world.

The components making up a tour, and in some cases even a cluster to be synchronised, can derive from other network environments - even though there may be no native support for such capabilities. The tour facility can also be used to seamlessly interface with other, non-networked media resources, such as material on a CD-ROM or hard drive somewhere else within the client's LAN.⁷ Thus large, regularly-used video or VR files can be held locally, then accessed as an integral part of globally networked interaction. (To appreciate the relevance of this, think about how long it would take to download 650 Mb of content through most existing Internet connections - or what the charges might be to do so!).




Figure 6

The architecture of the basic Hyper-G environment and its direct interoperation with other systems

The Hyper-G network environment is being co-developed by the Graz University of Technology in Austria and the University of Auckland in New Zealand, with about 140 person-years of work having gone into the project so far. The clients are being made freely available - as is the server software to educational and cultural organisations.⁸

Though the system is still in its infancy, it has already received considerable support. The European Space Agency, Austrian government, and Museum of New Zealand have adopted it for their own networks. All German universities and colleges will be setting up Hyper-G servers, as well as more than fifty universities world-wide that are installing a server by November '94 when the "Number 0" prototype edition of a premier electronic publication will appear. This new *Journal for Universal Computer Science* (J.UCS) is then scheduled to appear twelve times a year, beginning regular service in January of '95, as the first fully-refereed scientific publication to depend primarily upon Internet 'distribution'. It can thus provide the most current research reports and incorporate multimedia resources.⁹

As with Mosaic, however, attention tends to shift away from the underlying architecture and focus instead on the context of most people's actual experience - the user interface.

The Development of Harmony

The first Hyper-G client software has just been formally released, with the name *Harmony* selected for 'top of the line' viewers (i.e., ones which support all of the currently implemented capabilities of the environment). This is so far only available for systems running Unix X-Windows. A slightly more modest client is also available for MS-Windows, which is called *Amadeus* to distinguish it from the Harmony viewer. Work is in progress on a native Macintosh and PowerPC version of the viewer, with a first release scheduled for the end of '94.

The interactive design of these different clients is intended to preserve the look and feel of each platform environment, rather than impose a common interface style, but to still be similar enough in function that someone can readily feel 'at home' when moving from one to another. The overall intention is to constrain the contribution of 'authors' to optimised patterns of interaction - in order to allow the final user a personalised but consistent interface.¹⁰ A further project, presently in the planning stage, is the creation of a more refined 'public access terminal' version of Harmony for all those without a specific background of computer experience.

The X-Windows version of Harmony illustrates the range of different interface modes available. The *Session Manager* gives a hierarchical view of the material held in whatever server is currently being accessed. Selection is all point-and-click, with collections opening to show the layering of their contents and documents or clusters represented by icons that express their media constitution. The *Search* panel lets one select the nature, language, and target of a search routine . . . then displays the results as a weighted list. Downloading one of these updates the Session Manager, to show where that particular resource is located within the server providing it. The *Text Viewer* is similar to Mosaic's primary browser interface, with inline selectable images, hypertext links, and a graphic notation of past access paths (see next page, *Figure 7*).




Figure 7

The advantage of having links stored separately within the Hyper-G environment can be clearly seen on the *Local Map* of hypermedia relationships, which can be automatically generated for a selected document (Figure 8). This displays all of the outward references that are contained within that selection, along with those extending further onward from there within a chosen domain of network content, but it also shows any other documents within the same scope which contain references *to* the selected document. One can navigate using this hyperlink visualisation facility by picking another object of possible interest toward the edge of the map display, then generating another Local Map centered on that new selection. Remember that such mapping is ‘Local’ only in that it displays just a few levels of linkage relating to a particular selected document. The other documents which it shows could actually be on many different servers scattered all around the world.




Figure 8

Harmony's *History Browser* offers a timeline of one's past interactive waypoints, as another means of remaining oriented in hyperspace. Again, double-clicking will take you right back to any juncture, including a reconfigured Search panel. The *Image Viewer* provides various options for accessing a range of image formats, while its *Video* cousin is particularly fond of MPEG movies (*Figure 9*).

One telepresence application of these viewers is already being made available. Some server sites are beginning to supply very recent images or compressed video sequences which have been automatically acquired from a nearby camera, giving anyone anywhere on the Internet a chance to 'look in' from that viewpoint. Hyper-G's active protocol means that those doing such televideo can choose to be seeing only the most recently updated versions of the scene - and they can even remotely control the camera's direction and its zoom function in some cases.




Figure 9

A unique feature of Harmony is its support for three-dimensional world models (*Figure 10*). Users can manipulate 3-D scenes (translate, rotate, zoom), navigate within them (walk, fly, fly to), and interact with hypermedia links which can be expressed in a variety of different metaphors ('heads up' overlays, 'magic windows', embedded 'controls', etc.) - all within a networked context.¹¹ One obvious use for this capability is to provide what may be the only effective way to navigate within 'deep hyperspace' (very large-scale dynamic data structures with many tens of thousands of nodes): An interactive 'information *Landscape*'. This provides a 3-D visualisation of the contents within a section of the Hyper-G environment itself (*Figure 11*).




Figure 10




Figure 11

An exciting use for Harmony's 3-D abilities, and the one that brings together all of the other aspects of the Hyper-G network environment, is a kind of networked virtual reality. This will allow individuals from widespread localities to come together and meet each other in content-rich virtual worlds, either of the 'on-screen' or immersive variety. A telling precedent for this is the intense interest already being shown for the sort of mostly text-based virtual community called a *MOO* (MUD, Object Oriented). Participants can embed descriptions of themselves, create imagined 'rooms' in which to meet each other, and generally leave other descriptive or interactive artifacts in place.¹²

Once the support for this sort of networked 'meeting place' has been fully implemented within Hyper-G (and a rudimentary version is already available), the members of such communities will be able to actually see each other and move around within shared virtual domains, while at the same time having ready access to all of the hypermedia resources available on an exponentially expanding Internet.

A great deal of work remains to be done, but it is becoming easier to imagine a near future in which *anyone is able to be in contact with any other person or media resource, from almost anywhere, and is then able to easily share or do whatever can be managed through a digital network.*

References

- 1 Rheingold, H. **The Virtual Community**. Addison-Wesley, Reading, Mass., 1993.
- 2 Greenwald, J. *Wiring Africa*. **Wired** 2.06 , (June 1994), 60-62.
- 3 Myers, N. **The Gaia Atlas of Planet Management**. Doubleday, New York, 1985.
- 4 Frenkel, K. *A Conversation with Brenda Laurel*. **interactions** 1.1, (January 1994), 44-52.
- 5 Kappe, F., Maurer, H., and Sherbakov, N. *Hyper-G: A Universal Hypermedia System*. **Journal of Educational Multimedia and Hypermedia** 1, 2(1993), 39-66.
- 6 Nelson, T. **Literary Machines**. Distributors, South Bend, Indiana, 1987.
- 7 Andrews, K. and Kappe, F. *Soaring through hyperspace: a snapshot of Hyper-G and its Harmony client*. In **Multimedia/Hypermedia in Open Distributed Environments: Proceedings of EuroGraphics Symposium onMultimedia '94** (June 6-9,Graz, Austria), Springer-Verlag, New York, 1994, 181-191.
- 8 Information about *Hyper-G* and how to obtain client software is available from: **hyperg@cs.auckland.ac.nz** - in response to any e-mail.
- 9 Information about how to access the *Journal for Universal Computer Science* is available from: **jucs@iicm.tu-graz.ac.at** - Give [info] as 'Subject', or include "info" in text section.
- 10 Andrews, K. and Kappe, F. *Strait-Jacketing Authors: User Interface Consistency in Large-Scale Hypermedia Systems*. In **Proceedings of Hypermedia '93** (March, Zurich, Switzerland)
- 11 Andrews, K. *Constructing Cyberspace: Virtual Reality and Hypermedia*. In **Proceedings of Virtual Reality Vienna '93** (June 1993)
- 12 One of the most professional of these virtual communities, *Media MOO*, was established through the MIT Media Lab, and is serving as a virtual meeting place for those who are doing new media research. If you are able to visit online: **telnet purple-crayon.media.mit.edu 8888** then type "connect guest". Once connected you might start with "help" or "help introduction".